

Introduction to Neural Networks U. Minn. Psy 5038

More probability

■ Initialize standard library files:

```
In[7]:= Off[General::spell1];
```

The next package is needed for the add-on multivariate gaussian

```
In[8]:= << Statistics`MultinormalDistribution`
<< Statistics`DataManipulation`
```

```
In[10]:= << Graphics`Graphics`
```

Goals

Review the basics of probability distributions and statistics

More on generative modeling: drawing samples

Graphical models for inference

Optimal inference and Task dependence

Probability overview

Random variables, discrete probabilities, probability densities, cumulative distributions

■ Discrete: random variable X can take on a finite set of discrete values

$X = \{x(1), \dots, x(N)\}$

$$\sum_{i=1}^N p_i = \sum_{i=1}^N p(X = x(i)) = 1$$

■ Densities: X takes on continuous values, x, in some range.

Density: $p(x)$

Analogous to material mass,

we can think of the probability over some small domain of the random variable as "probability mass":

$$\begin{aligned} \text{prob}(x < X < dx + x) &= \int_x^{x+dx} p(x) dx \\ \text{prob}(x < X < dx + x) &\approx p(x) dx \end{aligned}$$

With the mass analogy, however, an object (event space) always "weighs" 1:

$$\int_{-\infty}^{\infty} p(x) dx = 1$$

Cumulative distribution:

$$\text{prob}(X < x) = \int_{-\infty}^x p(X) dX$$

■ Densities of discrete random variables

The Dirac Delta function, $\delta[\bullet]$, allows us to use the mathematics of continuous distributions for discrete ones, by defining the density as:

$$p[x] = \sum_{i=1}^N p_i \delta[x - x[i]], \text{ where } \delta[x - x[i]] = \begin{cases} \infty & \text{for } x = x[i] \\ 0 & \text{for } x \neq x[i] \end{cases}$$

Think of the delta function, $\delta[\bullet]$, as ϵ wide and $1/\epsilon$ tall, and then let $\epsilon \rightarrow 0$, so that:

$$\int_{-\infty}^{\infty} \delta(y) dy = 1$$

The density, $p[x]$, is a series of spikes. It is infinitely high only at those points for which $x = x[i]$, and zero elsewhere. But "infinity" is scaled so that the local mass or area around each point $x[i]$, is p_i .

■ Joint probabilities

$\text{Prob}(X \text{ AND } Y) = p(X, Y)$

Joint density: $p(x, y)$

Three basic rules of probability

Suppose we know everything there is to know about a set of variables (A,B,C,D,E). What does this mean in terms of probability? It means that we know the joint distribution, $p(A,B,C,D,E)$. In other words, for any particular combination of values ($A=a, B=b, C=c, D=d, E=e$), we can calculate, look up in a table, or determine some way or another the number $p(A=a, B=b, C=c, D=d, E=e)$.

Deterministic relationships are special cases.

■ Rule 1: Conditional probabilities from joints: The product rule

Probability about an event changes when new information is gained.

$\text{Prob}(X \text{ given } Y) = p(X|Y)$

$$p(X|Y) = \frac{p(X, Y)}{p(Y)}$$

$$p(X, Y) = p(X|Y) p(Y)$$

The form of the product rule is the same for densities as for probabilities.

■ Rule 2: Lower dimensional probabilities from joints: The sum rule (marginalization)

$$p(X) = \sum_{i=1}^N p(X, Y(i))$$

$$p(x) = \int_{-\infty}^{\infty} p(x, y) dx$$

■ Rule 3: Bayes' rule

From the product rule, and since $p[X,Y] = p[Y,X]$, we have:

$$p(Y|X) = \frac{p(X|Y)p(Y)}{p(X)}, \text{ and using the sum rule, } p(Y|X) = \frac{p(X|Y)p(Y)}{\sum_Y p(X, Y)}$$

■ Bayes Terminology in inference

Suppose we have some partial data (see half of someone's face), and we want to recall or complete the whole. Or suppose that we hear a voice, and from that visualize the face. These are both problems of statistical inference. We've already studied how to complete a partial pattern using energy minimization, and how energy minimization can be viewed as probability maximization.

We typically think of the **Y** term as a random variable over the hypothesis space (a face), and **X** as data or a stimulus (partial face, or sound). So for recalling a pattern **Y** from an input stimulus **X**, We'd like to have a function that tells us:

$p(\mathbf{Y} | \mathbf{X})$ which is called the **posterior** probability of the hypothesis (face) given the stimulus (partial face or sound).

-- i.e. what you get when you condition the joint by the stimulus data. The posterior is often what we'd like to base our decisions on, because it can be proved that picking the hypothesis **Y** which maximizes the posterior (i.e. maximum a posteriori or **MAP** estimation) minimizes the average probability of error.

$p(\mathbf{Y})$ is the **prior** probability of the hypothesis. Some hypotheses are more likely than others. Given a context, such as your room, some faces are a priori more likely than others. For me an image patch stimulating my retina in my kitchen is much more likely to be my wife's than my brother's (who lives in another state). This shows that priors are contingent, i.e. conditional on context, **p(Y| context)**.

$p(\mathbf{X}|\mathbf{Y})$ is the **likelihood** of the hypothesis. Note this is a probability of **X**, but not of **Y**. (The sum over **X** is one, but the sum over **Y** isn't necessarily one.)

■ Independence

Knowledge of one event doesn't change the probability of another event.

$$p(X)=p(X|Y)$$

$$p(X,Y)=p(X)p(Y)$$

Density mapping theorem

Suppose we have a change of variables that maps a discrete set of x 's uniquely to y 's: $X \rightarrow Y$.

■ Discrete random variables

No change to probability function. The mapping just corresponds to a change of labels, so the probabilities $p(X)=p(Y)$.

■ Continuous random variables

Form of probability density function does change because we require the probability "mass" to be unchanged: $p(x)dx = p(y)dy$

Suppose, $y=f(x)$

$$p_Y(Y) \delta Y = p_X(X) \delta X$$

In higher dimensions, the transformation is done by multiplying the density by the Jacobian, the determinant of the matrix of partial derivatives of the change of coordinates.

$$p_Y(y) = \int \delta(y - f(x)) f^{-1}(x) p_X(x) dx$$

over each monotonic part of f.

Convolution theorem for adding rvs

Let x be distributed as g(x), and y as h(x). Then the probability density for z=x+y is, f(z):

$$f(z) = \int g(s) h(z-s) ds \quad (1)$$

Statistics

■ Expectation & variance

Analogous to center of mass:

Definition of expectation or average:

$$\text{Average}[X] = \bar{X} = E[X] = \sum x[i] p[x[i]] \sim \sum_{i=1}^N x_i / N$$

$$\mu = E[X] = \int x p(x) dx$$

Some rules:

$$E[X+Y] = E[X] + E[Y]$$

$$E[aX] = aE[X]$$

$$E[X+a] = a + E[X]$$

Definition of variance:

$$\sigma^2 = \text{Var}[X] = E[(X-\mu)^2] = \sum_{j=1}^N (p(x(j))) (x(j) - \mu)^2 = \sum_{j=1}^N p_j (x_j - \mu)^2$$

$$\text{Var}[X] = \int (x - \mu)^2 p(x) dx \sim \sum_{i=1}^N (x_i - \mu)^2 / N$$

Standard deviation:

$$\sigma = \sqrt{\text{Var}[X]}$$

Some rules:

$$\begin{aligned} \text{Var}[X] &= E[X^2] - E[X]^2 \\ \text{Var}[aX] &= a^2 \text{Var}[X] \end{aligned}$$

■ Covariance & Correlation

Covariance:

$$\text{Cov}[X, Y] = E[(X - \mu_X)(Y - \mu_Y)]$$

Correlation coefficient:

$$\rho[X, Y] = \frac{\text{Cov}[X, Y]}{\sigma_X \sigma_Y}$$

■ Cross and Autocovariance matrix

Suppose X and Y are vectors: $\{X_1, X_2, \dots\}$ and $\{Y_1, Y_2, \dots\}$

$$\text{Cov}[X_i, Y_j] = E[(X_i - \mu_{X_i})(Y_j - \mu_{Y_j})] \sim \sum_{n=1}^N (x_i^n - \mu_{X_i})(y_j^n - \mu_{Y_j}) / N$$

$$\text{Autocov}[X_i, X_j] = E[(X_i - \mu_{X_i})(X_j - \mu_{X_j})] \sim \sum_{n=1}^N (x_i^n - \mu_{X_i})(x_j^n - \mu_{X_j}) / N$$

In other words, the autocovariance matrix can be approximated by the average outer product. It is a Hebbian matrix memory of pair-wise relationships.

■ Independent random variables

If $p(X, Y) = p(X)p(Y)$, then

$$E[XY] = E[X]E[Y] \text{ (uncorrelated)}$$

$$\text{Cov}[X, Y] = \rho[X, Y] = 0$$

$$\text{Var}[X+Y] = \text{Var}[X] + \text{Var}[Y]$$

If two random variables are uncorrelated, they are not necessarily independent.

Two random variables are said to be orthogonal if their correlation is zero.

Degree of belief vs., relative frequency

What is the probability that the Vikings will win the Superbowl in 2004? Assigning a number between 0 and 1 is assigning a degree of belief. These probabilities are also called subjective probabilities. "Odds" determine subjective probabilities, where the "odds of x to y" means probability = $x/(x+y)$.

What is the probability that a coin will come up heads? In this case, we can do an experiment. Flip the coin n times, and count the number of heads, say h[n], and then set the probability, $p = h[n]/n$ -- the relative frequency. Of course, if we did it again, we may not get the same estimate of p. One solution often given is:

$$p = \lim_{n \rightarrow \infty} \frac{h(n)}{n}$$

A problem with this, is that there is no guarantee that a well – defined limit exists.

In some domains we can measure statistics, and model probabilities of both inputs and outputs. So the relative frequency interpretation seems reasonable. In practice, the dimensions of many problems in perception, cognition, and memory are so high, that it is impractical to do this. Once we use the statistical framework to model perception, say of a particular cue (say), then probabilities are more like "subjective unconscious beliefs".

Principle of insufficient reason

■ Principle of symmetry

Suppose we have N events, $x[1], x[2], x[3], \dots, x[N]$ that are all physically identical except for the label. Then assume that

$$\text{prob}(x(1)) = \text{prob}(x(2)) = \text{prob}(x(3)) = \text{prob}(x(N)) = \frac{1}{N}$$

In other words, if we have no additional information about the events, we should assume that they are uniformly distributed. I.e., assume a uniform prior.

What about the continuous case where there is no reason to assume any particular value at all between $-\infty$ and $+\infty$?

Improper priors.

■ Information theory and Maximum entropy

Information theory provides a powerful extension to the principle of symmetry. Information of event X is:

$$\text{Information}[X] = -\log_2(p(X))$$

Using the definition of expectation above, we can specify the expectation of information, which is called entropy. Entropy of a random variable X with probability distribution $p[X]$ is:

$$H(X) = \text{Average}(\text{Information}[X]) = -\sum_X p(X) \log_2(p(X))$$

It can be shown that out of all possible probability distributions, $H(X)$ is biggest for the uniform distribution, $p(X)=1/N$. Maximum entropy is looking like symmetry.

It turns out that a more powerful formulation of the principle of symmetry is maximum entropy. For example, out of all possible probability distributions of a random variable with infinity range, but with a specific mean and standard deviation, the Gaussian is unique in having the largest entropy. If the range goes from zero to infinity, and we know the mean, the maximum entropy distribution is an exponential (Cover and Thomas).

An interesting application of the maximum entropy principle is to learning image textures joint probabilities: $p(I[1], \dots, I[N])$, where N is very big, but where one has only a relatively small number of measured statistics relative to the number of possible images (which is really huge). The measurements underdetermine the dimensionality of the probability space – i.e. there are many different probability distributions which give the same statistics. So the principle of symmetry, or insufficient reason, says to choose the one with the maximum entropy.


More on generative modeling: Multivariate gaussian, mixtures

Making a univariate (scalar) gaussian random number generator: Method 1: Just for Gaussian

■ Use Central Limit Theorem

If all we want to do is make a Gaussian random number generator from a uniformly distributed generator, we can use the Central Limit Theorem.

```
nusamples = 10;
z1 = Table[Sum[Random[], {i, nusamples}] - nusamples / 2, {1000}];
binsize = .1;
midpoints = Table[i + binsize / 2, {i, -3, 3 - binsize, binsize}];
freq = BinCounts[z1, {-3, 3, binsize}];
BarChart[Transpose[{freq, midpoints}], BarLabels -> None];
```



Method 2: Use Density Mapping theorem. More general.

We'll use the density mapping theorem to turn uniformly distributed random numbers **Random[]** into gaussian distributed random numbers with mean =0 and standard deviation =1.

$$p_Y(Y) \delta Y = p_X(X) \delta X$$

$$p_Y (Y) \frac{\delta Y}{\delta X} = p_X (X)$$

Suppose $p_Y (Y) = 1$ (over the unit interval, but zero elsewhere). Then

$$Y (X) = \int_{-\infty}^X p_X (x') dx' = P (X) \quad (2)$$

Thus if we sample from the uniform distribution to get y , x should be distributed according to $p_X (X)$. To do this, we need a mapping from $y \rightarrow x$. This is given by the inverse cumulative distribution, i.e. $P^{-1}(y)$.

Let's implement this. The quick way is to use *Mathematica's* built-in function to get the inverse cumulative.


Method 2a: Applied to Gaussian

`InverseErf[]` is the inverse of :


$$\text{erf} (z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$$

We can use this to define a function for the inverse cumulative of a gaussian:

```
z[p_] := Sqrt[2] InverseErf[1 - 2 p];
Plot[z[y], {y, 0, 1}];
```



```
binsize = .1;
midpoints = Table[i + binsize/2, {i, -3, 3 - binsize, binsize}];
z1 = Table[z[Random[]], {5000}];
freq = BinCounts[z1, {-3, 3, binsize}];
BarChart[Transpose[{freq, midpoints}], BarLabels -> None];
```



Method 2b: From scratch: Works for almost any distribution. Do the Gaussian case for a third time


Suppose we have a discrete representation of any cumulative distribution. How can we generate samples? We'll illustrate the method with a discretization of the Gaussian. Our immediate goal is to produce a discrete approximation to the cumulative gaussian. To review where things come from, we'll start with the definition of a Gaussian, and make sure it is normalized.

```
Integrate[Exp[-(x - x0)^2 / (2 * sigma^2)], {x, -Infinity, Infinity}]
```

```
If[Re[sigma^2] > 0, Sqrt[2 * Pi] Sqrt[sigma^2], Integrate[Exp[-(x-x0)^2 / (2 * sigma^2)], {x, -Infinity, Infinity}]]
```

Let $x_0=0$ and $\sigma=1$:

```
Plot[Exp[-(x1^2) / 2] / (Sqrt[2 * Pi]), {x1, -4, 4}];
```



`Plot[PDF[NormalDistribution[0,1],x1],{x1,-4,4}];` gives the same thing using the add-on normal distribution function.


■ Cumulative gaussian

```
Clear[cumulgauss, x, x1];
cumulgauss[x_] := Integrate[Exp[-(x1^2) / 2] / (Sqrt[2 * Pi]), {x1, -Infinity, x}];
```


```
cumulgauss[Infinity]
```

```
1
```

```
Plot[cumulgauss[x], {x, -4, 4}];
```



```
lcumulgauss = Table[{x, cumulgauss[x]}, {x, -4.0, 4.0, .25}];
ListPlot[lcumulgauss];
```




■ Make inverse cumulative gaussian table

```
invcumulgauss = RotateLeft[lcumulgauss, {1, 1}];
```

To see what this does, evaluate:

```
In[5]:= {{x1, y1}, {x2, y2}, {x3, y3}}
RotateLeft[{{x1, y1}, {x2, y2}, {x3, y3}}, {1, 1}]
```


```
ListPlot[invcumulgauss];
```



■ Make interpolated function of the inverse cumulative

```
interinvcumulgauss = Interpolation[invcumulgauss];
```

```
Plot[interinvcumulgauss[x], {x, .01, .99}];
```



■ Draw samples with a standard deviation of Sqrt[10]

```
Round[10 * interinvcumulgauss[Random[]]]
```

```
-1
```


■ Draw a bunch of samples, and plot up histogram

```
z = Table[Round[10 * interinvcumulgauss[Random[]]], {10000}];
domain = Range[-20, 20];
Freq = Map[Count[z, #] &, domain];
```


```
Sqrt[10.]
```

```
3.16228
```

```
ListPlot[Freq];
```




```
<< Statistics`DataSmoothing`
ListPlot[MovingAverage[Freq, 6]];
```



■ Plot up cumulative histogram

```
CumFreq = FoldList[Plus, 0, Freq];
ListPlot[CumFreq];
```



```
Transpose[{domain, Freq, CumFreq}] // MatrixForm
```

Multivariate (vector) gaussian distributions

■ Define multivariate gaussian probability density

An n -variate multivariate gaussian (multinormal) distribution with mean vector μ and covariance matrix Σ is denoted $N_n(\mu, \Sigma)$. The density is:

$$p(x) = \frac{1}{(2\pi)^{n/2} \text{Det}[\Sigma]^{1/2}} \text{Exp}\left[-\frac{1}{2} (x - \mu)^T \Sigma^{-1} (x - \mu)\right] \quad (3)$$


We define a 2-variate density:

```
multigaus[x_, m_, cov_] :=
Module[{IC, detCov, norm, p},
IC = Inverse[cov];
detCov = Abs[Det[cov]];
norm = N[Sqrt[(2Pi)^2 detCov]];
p = Exp[-0.5 (x-m).IC.(x-m)]/norm;
Return[p];
];
```

■ Two variable examples


■ Zero mean, zero correlation

```
m1 = {0,0};
Cov = {{1,0},{0,1}};
ContourPlot[multigauss[{x1,x2},m1,Cov],{x1,-2,2}, {x2,-2,2}];
```




■ Mean = {1,1}, positive correlation

```
m1 = {1,1};
Cov = {{1,.5},{.5,1}};
ContourPlot[multigauss[{x1,x2},m1,Cov],{x1,0,2}, {x2,0,2}];
```



■ Mean = {1,1}, positive correlation, small variance

```
m1 = {1,1};
Cov = 0.2{{1,.5},{.5,1}};
ContourPlot[multigauss[{x1,x2},m1,Cov],{x1,0,2}, {x2,0,2}];
```



■ Mixture of gaussians

α is a mixing parameter

$$p(x) = \alpha p_1(x) + (1 - \alpha) p_2(x) \text{ where } 0 \leq \alpha \leq 1 \quad (4)$$

α can be interpreted in terms of a prior probability of choosing which of two distributions a sample will be drawn from.


Starting with $p(x,a)$ use the sum and product rules to derive the above mixture density where

$p(a=a_1) = \alpha$, and $p(a=a_2) = (1-\alpha)$.

```
m1 = {1,.5}; m2 = {-1,-.5};
Cov1 = 0.4*{{1,.6},{.6,1}};
Cov2 = 0.4*{{1,-.6},{-.6,1}};
mix[x_] := 0.5 (multigauss[x,m1,Cov1] + multigauss[x,m2,Cov2]);
```



```
ContourPlot[mix[{x1,x2}], {x1,-2,2}, {x2,-2,2}];
```



■ Drawing samples from the density--draw from a hat method

- We'll simulate the process of filling a hat with slips of paper, where the number of slips is proportional to the probability the number being in some range (dx1,dx2)

```
m1 = {0,0};
Cov = {{1,.8},{.8,1}};
dx1 = 0.1;
dx2 = dx1;

Nslips=100;
hat = {};
For[x1=-2,x1<=2,x1=x1+dx1,
  For[x2=-2,x2<=2,x2=x2+dx2,
    np = Nslips*multigauss[{x1,x2},m1,Cov];
    For[i=1,i<np,i=i+1,
      hat = Append[hat,{x1,x2}];
    ];
  ];
];
```

hat is a list of pairs of numbers for which the frequency of occurrence of pairs is determined by multigauss.

```
Dimensions[hat]
```

```
{8698, 2}
```

■ Now let's do a check, where we compile a histogram representing the frequencies of each slip

First, define the "bins" in domain, that we'll use to check for matches:

```
domain = {};
For[x1=-2,x1<=2,x1=x1+dx1,
  For[x2=-2,x2<=2,x2=x2+dx2,
    domain = Append[domain,{x1,x2}];
  ];
];
```

An alternate way of specifying the domain using Outer[], and Range[]:

```
domain2 = Flatten[Outer[List,Range[-2,2,dx1],Range[-2,2,dx1]],1];
```


Now we'll count how many times we find that an element of hat matches a domain element:

```
Freq = Map[Count[hat,#]&,domain];
```

```
width = Sqrt[Dimensions[Freq]]
```

```
{40}
```

```
ListDensityPlot[Partition[Freq,width]];
```




■ Now draw a sample--simulate pulling a slip from hat

```
rv:=hat[[Random[Integer,{1,Length[hat]}]]];
```

```
test = Table[hat[[Random[Integer,{1,Length[hat]}]]],{600}];
```


```
g1=ListPlot[test];
```



Of course, we can't see the frequency of draws in this plot, so let's count up the number of occurrences per bin, and plot up the results as we did above.

```
Freq2 = Map[Count[test, #]&, domain];
width = Sqrt[Dimensions[Freq2]];
```

```
ListDensityPlot[Partition[Freq2, width]];
```



Exercise: Drawing multivariate samples from the density -- use the inverse cumulative distribution

Add-on *Mathematica* functions for gaussian multivariates & exploring marginals

■ Define PDF, CDF

The add-on function, `MultinormalDistribution[mu, sigma]` specifies a multinormal (multivariate Gaussian) distribution with mean

```
m1 = {1, .5};
r=0.4*{{1, .6}, {.6, 4}};
ndist = MultinormalDistribution[m1, r];
```

```
pdf = PDF[ndist, {x1, x2}]
```


```
0.20855
e1/2 (-(-1+x1) (2.74725 (-1+x1)-0.412088 (-0.5+x2))-(-0.412088 (-1+x1)+0.686813 (-0.5+x2)) (-0.5+x2))
```

What is the probability of the distribution in the region $x_1 < -2 \cap x_2 < 1$.


```
CDF[ndist, {-2, 1}]
```

```
1.02471 × 10-6
```


```
g1=ContourPlot[PDF[ndist, {x1, x2}], {x1, -2, 2}, {x2, -2, 2}, ContourShading->False];
```



```
marginal[x1_] := Integrate[PDF[ndist, {x1, x2}], {x2, -Infinity, Infinity}];
g2 = Plot[marginal[x1], {x1, -2, 2}];
```



```
Show[{g1, g2}];
```



■ Drawing samples

As we've used in earlier lectures, drawing samples is done by:

```
Random[ndist]
```

```
{-0.52553, -2.54926}
```


■ Mixtures of gaussians

```
Clear[mix];
```

```
r1=0.4*{1,.6},{.6,1};
r2=0.4*{1,-.6},{-.6,1};
m1 = {1,.5}; m2 = {-1,-.5};
ndist1 = MultinormalDistribution[m1, r1];
ndist2 = MultinormalDistribution[m2, r2];
```

```
mix[x_] := 0.5 (PDF[ndist1, x] + PDF[ndist2, x]);
```

```
ContourPlot[mix[{x1,x2}],{x1,-2,2}, {x2,-2,2}];
```




■ Marginals for mixture

```
marginal[x1_] := Integrate[mix[{x1, x2}], {x2, -Infinity, Infinity}] (5)
```


```
Clear[marginal];
marginal[x1_] := 0.5 * (Integrate[PDF[ndist1, {x1, x2}], {x2, -Infinity, Infinity}] +
Integrate[PDF[ndist2, {x1, x2}], {x2, -Infinity, Infinity}]);
```

```
Plot[marginal[x1], {x1, -2, 2}];
```



```
Clear[marginal];
marginal[x2_] := 0.5 * (Integrate[PDF[ndist1, {x1, x2}], {x1, -Infinity, Infinity}] +
Integrate[PDF[ndist2, {x1, x2}], {x1, -Infinity, Infinity}]);
```

```
Plot[marginal[x2], {x2, -2, 2}];
```




Which projection (marginal) is more "interesting"—the one onto x1 or onto x2?
Exploratory projection pursuit

Graphical Models of dependence

■ Graphs: causal structure and conditional independence

The idea is to represent the probabilistic structure of the joint distribution $P(S,L,I)$ by a Bayes net (e.g. Ripley, 1996), which is a graphical model that expresses how variables influence each other. There are just three basic building blocks: converging, diverging, and intermediate nodes. For example, multiple causal variables causing a given measurement, a single variable producing multiple measurements, or a cause indirectly influencing a measurement through an intermediate variable. These types of influence provide a first step towards modeling the joint distribution and the means to compute probabilities of the unknown variables given known values.



Components of the generative structure for data patterns involve converging, diverging, and intermediate nodes. For example, these could correspond to: multiple (scene) causes {shape S_1 , illumination S_2 } giving rise to the same image measurement, I ; one cause, S influencing more than one image measurement, {color, I_1 , brightness, I_2 }; a scene (or other) cause S , {object identity, S } influencing an image measurement (image contour) through an intermediate variable L (3D shape).


The arrows tell us how to factor the joint probability into conditionals. So for the three examples above, we have:

$$p(S_1, S_2, I) = p(I|S_1, S_2)p(S_1)p(S_2)$$

$$p(S_1, I_1, I_2) = p(I_1|S)p(I_2|S)p(S)$$

$$p(S, L, I) = p(I|L)p(L|S)p(S)$$

Graphical models can provide a way of sketching out the conditional relationships between a complex set of interactions. Below is an example (see Kersten, Mamassian and Yuille, 2003) of graphical model for high-level vision.



■ Primary, secondary variables.

We can interpret the causal structure in terms of conditional probability.

The data measurements (x) are determined by a typically non-linear function (ϕ) of primary signal variables (S_e) (to be explicitly estimated) and confounding secondary variables (S_g) (also called "generic" variables to be discounted). Knowledge is represented by the joint probability $p(x, S_e, S_g)$. In general, the causal structure of natural data (e.g. image or speech) patterns is more complex and consequently requires elaboration of its graphical representation. For pattern inference theory, the task is to make a decision about the signal hypotheses or primary signal variables, while discounting the noise or secondary variables. Thus optimal perceptual decisions are determined by $p(x, S_e)$, which is derived by summing over the secondary variables (i.e. marginalizing with respect to the secondary variables): $\int_{S_g} p(x, S_e, S_g) dS_g$.

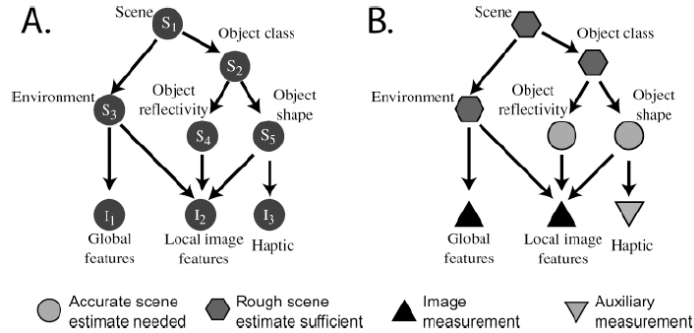
Influences between variables are represented by conditioning, and a graphical model expresses the conditional independencies between variables. Two random variables may only become independent, however, once the value of some third variable is known. This is called conditional independence. Recall from above that two random variables are independent if and only if their joint probability is equal to the product of their individual probabilities. Thus, if $p(A, B) = p(A)p(B)$, then A and B are independent. If $p(A, B|C) = p(A|C)p(B|C)$, then A and B are conditionally independent.

When corn prices drop in the summer, hay fever incidence goes up. However, if the joint on corn price and hay fever is conditioned on "ideal weather for corn and ragweed", the correlation between corn prices and hay fever drops. This is because corn price and hay fever symptoms are conditionally independent.

There is a correlation between eating ice cream and drowning. Why? What event should you condition on to make the dependence go away?

■ What is noise? Primary and secondary variables

The primary variables need to be estimated accurately. Noise is whatever you don't care to estimate, but contributes to the data. The secondary variables are noise, where either no estimate is required, or perhaps only a rough estimate.



Optimal Inference and task dependence: Fruit example

(due to James Coughlan; see Yuille, Coughlan, Kersten & Schrater).




Figure from Yuille, Coughlan, Kersten & Schrater.

The graph specifies how to decompose the joint probability:

$$p[F, C, I_S, I_C] = p[I_C | C] p[C | F] p[I_S | F] p[F]$$

The prior model on hypotheses, F & C

More apples (F=1) than tomatoes (F=2), and:

```
ppF[F_] := If[F == 1, 9 / 16, 7 / 16];
TableForm[Table[ppF[F], {F, 1, 2}], TableHeadings -> {"F=a", "F=t"}]
```

F=a	$\frac{9}{16}$
F=t	$\frac{7}{16}$

The conditional probability cpCF[C|F]:

```
cpCF[F_, C_] := Which[F == 1 && C == 1, 5 / 9,
    F == 1 && C == 2, 4 / 9, F == 2 && C == 1, 6 / 7, F == 2 && C == 2, 1 / 7];
TableForm[Table[cpCF[F, C], {F, 1, 2}, {C, 1, 2}],
    TableHeadings -> {"F=a", "F=t"}, {"C=r", "C=g"}]
```

	C=r	C=g
F=a	$\frac{5}{9}$	$\frac{4}{9}$
F=t	$\frac{6}{7}$	$\frac{1}{7}$

So the joint is:

```
jpFC[F_, C_] := cpCF[F, C] ppF[F];
TableForm[Table[jpFC[F, C], {F, 1, 2}, {C, 1, 2}],
    TableHeadings -> {"F=a", "F=t"}, {"C=r", "C=g"}]
```

	C=r	C=g
F=a	$\frac{5}{16}$	$\frac{1}{4}$
F=t	$\frac{3}{8}$	$\frac{1}{16}$

We can marginalize to get the prior probability on color alone is:

$$ppC[C_] := \sum_{F=1}^2 jpFC[F, C]$$

Question: Is fruit identity independent of material color--i.e. is F independent of C?

■ Answer

No.

```
TableForm[Table[jpFC[F, C], {F, 1, 2}, {C, 1, 2}],
  TableHeadings -> {{ "F=a", "F=t" }, { "C=r", "C=g" }}]
TableForm[Table[ppF[F] ppC[C], {F, 1, 2}, {C, 1, 2}],
  TableHeadings -> {{ "F=a", "F=t" }, { "C=r", "C=g" }}]
```

	C=r	C=g
F=a	$\frac{5}{16}$	$\frac{1}{4}$
F=t	$\frac{3}{8}$	$\frac{1}{16}$

	C=r	C=g
F=a	$\frac{99}{256}$	$\frac{45}{256}$
F=t	$\frac{77}{256}$	$\frac{35}{256}$

The generative model: Imaging probabilities

Suppose that we have gathered some "image statistics" which provides us knowledge of how the image measurements for shape Is , and for color Ic depend on the type of fruit F , and material color, C . For simplicity, our measurements are discrete and binary (a more realistic case, they would have continuous values), say $Is = \{am, tm\}$, and $Ic = \{rm, gm\}$.

$$P(I_S=am, tm | F=a) = \{11/16, 5/16\}$$

$$P(I_S=am, tm | F=t) = \{5/8, 3/8\}$$

$$P(I_C=rm, gm | C=r) = \{9/16, 7/16\}$$

$$P(I_C=rm, gm | C=g) = \{1/2, 1/2\}$$

We use the notation am , tm , rm , gm because the measurements are already suggestive of the likely cause. So there is a correlation between apple and apple-like shapes, am ; and between red material, and "red" measurements. In general, there may not be an obvious correlation like this.

We define a function for the probability of Ic given C , $cpIcC[Ic | C]$:

```
cpIcC[Ic_, C_] := Which[Ic == 1 && C == 1, 9/16,
  Ic == 1 && C == 2, 7/16, Ic == 2 && C == 1, 1/2, Ic == 2 && C == 2, 1/2];
TableForm[Table[cpIcC[Ic, C], {Ic, 1, 2}, {C, 1, 2}],
  TableHeadings -> {{ "Ic=rm", "Ic=gm" }, { "C=r", "C=g" }}]
```

	C=r	C=g
Ic=rm	$\frac{9}{16}$	$\frac{7}{16}$
Ic=gm	$\frac{1}{2}$	$\frac{1}{2}$

The probability of Is conditional on F is $cpIsF[Is | F]$:

```
cpIsF[Is_, F_] := Which[Is == 1 && F == 1, 11/16,
  Is == 1 && F == 2, 5/8, Is == 2 && F == 1, 5/16, Is == 2 && F == 2, 3/8];
TableForm[Table[cpIsF[Is, F], {Is, 1, 2}, {F, 1, 2}],
  TableHeadings -> {{ "Is=am", "Is=tm" }, { "F=a", "F=t" }}]
```

	F=a	F=t
Is=am	$\frac{11}{16}$	$\frac{5}{8}$
Is=tm	$\frac{5}{16}$	$\frac{3}{8}$

The total joint probability

We now have enough information to put probabilities on the $2 \times 2 \times 2$ "universe" of possibilities, i.e. all possible combinations of fruit, color, and image measurements. Looking at the graphical model makes it easy to use the product rule to construct the total joint, which is:

$$p[F, C, Is, Ic] = p[Ic | C] p[C | F] p[Is | F] p[F]$$

```
jpFCIsIc[F_, C_, Is_, Ic_] := cpIcC[Ic, C] cpCF[F, C] cpIsF[Is, F] ppF[F]
```

Usually, we don't need the probabilities of the image measurements (because once the measurements are made, they are fixed and we want to compare the probabilities of the hypotheses. But in our simple case here, once we have the joint, we can calculate the probabilities of the image measurements through marginalization $p(Is, Ic) = \sum_C \sum_F p(F, C, Is, Ic)$, too:

```
jpIsIc[Is_, Ic_] := Sum[Sum[jpFCIsIc[F, C, Is, Ic], {C, 1, 2}], {F, 1, 2}]
```

Three MAP tasks

We are going to show that the best guess (i.e. maximum probability) depends on the task.

■ Define argmax[] function:

```
argmax[x_] := Position[x, Max[x]];
```

■ Pick most probable fruit AND color--Answer "red tomato"

First, suppose the task is to make the best bet as to the fruit AND material color. To make it concrete, suppose that we see an "apple-like shape" with a reddish color, i.e., we measure $I_s=am=1$, and $I_c=rm=1$. The measurements suggest "red apple", but to find the most probable, we need to take into account the priors too in order to make the best guesses.

Using the total joint, $p(F,C | I_s, I_c) = \frac{p(F,C,I_s,I_c)}{p(I_s,I_c)} \propto p(F,C,I_s=1, I_c=1)$

```
TableForm[jpFCIsIcTable = Table[jpFCIsIc[F, C, 1, 1], {F, 1, 2}, {C, 1, 2}],
TableHeadings -> {"F=a", "F=t"}, {"C=r", "C=g"}]]
Max[jpFCIsIcTable]
argmax[jpFCIsIcTable]
```

	C=r	C=g
F=a	$\frac{495}{4096}$	$\frac{77}{1024}$
F=t	$\frac{135}{1024}$	$\frac{35}{2048}$

$\frac{135}{1024}$

{{2, 1}}

"Red tomato" is the most probable once we take into account the difference in priors.

Calculating $p(F,C | I_s, I_c)$. We didn't actually need $p(F,C | I_s, I_c)$, but we can calculate it by conditioning the total joint on the probability of the measurements:

```
jpFCcIsIc[F_, C_, Is_, Ic_] := jpFCIsIc[F, C, Is, Ic] / jpIsIc[Is, Ic]
```

```
TableForm[jpFCcIsIcTable = Table[jpFCcIsIc[F, C, 1, 1], {F, 1, 2}, {C, 1, 2}],
TableHeadings -> {"F=a", "F=t"}, {"C=r", "C=g"}]]
Max[jpFCcIsIcTable]
argmax[jpFCcIsIcTable]
```

	C=r	C=g
F=a	$\frac{55}{157}$	$\frac{308}{1413}$
F=t	$\frac{60}{157}$	$\frac{70}{1413}$

$\frac{60}{157}$

{{2, 1}}

■ Pick most probable color--Answer "red"

Same measurements as before. But now suppose we only care about the true material color, and not the identity of the object. Then we want to integrate out or marginalize with respect to the shape variable, F. In this case, we want to maximize the posterior:

$$p(C | I_s=1, I_c=1) = \sum_{F=1}^2 p(F, C | I_s=1, I_c=1)$$

```
pC[C_, Is_, Ic_] := Sum[jpFCcIsIc[F, C, Is, Ic], {F, 1, 2}]
```

```
TableForm[pCTable = Table[pC[C, 1, 1], {C, 1, 2}], TableHeadings -> {"C=r", "C=g"}]]
Max[pCTable]
argmax[pCTable]
```

C=r	$\frac{115}{157}$
C=g	$\frac{42}{157}$

$\frac{115}{157}$

{{1}}

Answer is that the most probable material color is C = r, "red".

■ Pick most probable fruit--Answer "apple"

Same measurements as before. But now, we don't care about the material color, just the identity of the fruit. Then we want to integrate out or marginalize with respect to the material variable, C. In this case, we want to maximize the posterior:

$p(F | I_s, I_c)$

$$pF[F_, Is_, Ic_] := \sum_{C=1}^2 j pFCcIsIc[F, C, Is, Ic]$$

```
TableForm[pFTable = Table[pF[F, 1, 1], {F, 1, 2}], TableHeadings -> {"F=a", "F=t"}]
Max[pFTable]
argmax[pFTable]
```

F=a	$\frac{803}{1413}$
F=t	$\frac{610}{1413}$

$$\frac{803}{1413}$$

```
{{1}}
```

The answer is "apple". So to sum up, for the same data measurements, the most probable fruit AND color is "red tomato", but the most probable fruit is "apple"!

■ Important "take-home message": *Optimal inference depends on the precise definition of the task*

Try expressing the consequences using the frequency interpretation of probability.

Appendices

```
<< Graphics`Graphics`
```

Using *Mathematica* lists to manipulate discrete priors, likelihoods, and posteriors

■ A note on list arithmetic

We haven't done standard matrix/vector operations above to do conditioning. We've take advantage of how *Mathematica* divides a 2x3 array by a 2-element vector:

```
M=Array[m,{2,3}]
x = Array[x,{2}]
```

$$\begin{pmatrix} m(1,1) & m(1,2) & m(1,3) \\ m(2,1) & m(2,2) & m(2,3) \end{pmatrix}$$

$$\{x(1), x(2)\}$$

```
M/x
```

$$\begin{pmatrix} \frac{m(1,1)}{x(1)} & \frac{m(1,2)}{x(1)} & \frac{m(1,3)}{x(1)} \\ \frac{m(2,1)}{x(2)} & \frac{m(2,2)}{x(2)} & \frac{m(2,3)}{x(2)} \end{pmatrix}$$

■ Putting the probabilities back together again to get the joint

```
Transpose[Transpose[pHx] px]
```

$$\begin{pmatrix} \frac{1}{12} & \frac{1}{12} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{6} & \frac{1}{6} \end{pmatrix}$$

```
pxH pH
```

$$\begin{pmatrix} \frac{1}{12} & \frac{1}{12} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{6} & \frac{1}{6} \end{pmatrix}$$

■ Getting the posterior from the priors and likelihoods:

One reason Bayes' theorem is so useful is that it is often easier to formulate the likelihoods (e.g. from a causal or generative-model of how the data could have occurred), and the priors (often from heuristics, or in computational vision empirically testable models of the external visual world). So let's use *Mathematica* to derive $p(\mathbf{H}|\mathbf{x})$ from $p(\mathbf{x}|\mathbf{H})$ and $p(\mathbf{H})$. (i.e. $p(\mathbf{H}|\mathbf{x})$ from $p(\mathbf{x}|\mathbf{H})$ and $p(\mathbf{H})$).

```
px2 = Plus @@ (pxH pH)
```

```
{ $\frac{5}{12}, \frac{1}{4}, \frac{1}{3}$ }
```

```
Transpose[Transpose[pxH pH] / Plus @@ (pxH pH)]
```

```
{ $\frac{1}{5}, \frac{1}{3}, \frac{1}{2}$ ;  
 $\frac{4}{5}, \frac{2}{3}, \frac{1}{2}$ }
```

■ Show that this joint probability has a uniform prior (i.e. both priors equal).

```
p = {{1/8, 1/8, 1/4}, {1/4, 1/8, 1/8}}
```

```
{{ $\frac{1}{8}, \frac{1}{8}, \frac{1}{4}$ }, { $\frac{1}{4}, \frac{1}{8}, \frac{1}{8}$ }}
```

Marginalization and conditioning: A small dimensional example using list manipulation in *Mathematica*

■ A discrete joint probability

All of our knowledge regarding the signal discrimination problem can be described in terms of the joint probability of the hypotheses, \mathbf{H} and the possible data measurements, \mathbf{x} . The probability function assigns a number to all possible combinations:

$p(\mathbf{H}, \mathbf{x})$

That is, we are assuming that both the hypotheses and the data are discrete random variables.

```
H = {S1  
S2  
x ∈ {1, 2, ...}}
```

Let's assume that x can only take on one of three values, 1, 2, or 3. And suppose the joint probability is:

```
P = {{ $\frac{1}{12}, \frac{1}{12}, \frac{1}{6}$ }, { $\frac{1}{3}, \frac{1}{6}, \frac{1}{6}$ }}
```

```
{ $\frac{1}{12}, \frac{1}{12}, \frac{1}{6}$ ;  
 $\frac{1}{3}, \frac{1}{6}, \frac{1}{6}$ }
```

```
TableForm[p, TableHeadings -> {"H=S1", "H=S2"}, {"x=1", "x=2", "x=3"}]
```

	x=1	x=2	x=3
H=S1	$\frac{1}{12}$	$\frac{1}{12}$	$\frac{1}{6}$
H=S2	$\frac{1}{3}$	$\frac{1}{6}$	$\frac{1}{6}$

The total probability should sum up to one. Let's test to make sure. We first turn the list of lists into a single list of scalars using `Flatten[]`. And then we can sum either with `Apply[Plus, Flatten[p]]`.

```
Plus @@ Flatten[p]
```

```
1
```

We can pull out the first row of p like this:

```
p[[1]]
```

```
{ $\frac{1}{12}, \frac{1}{12}, \frac{1}{6}$ }
```

Is this the probability of x ? No. For a start, the numbers don't sum to one. But we can get it through the two processes of marginalization and conditioning.

■ Marginalizing

What are the probabilities of the data, $p(x)$? To find out, we use the *sum rule* to sum over the columns:

```
pX = Apply[Plus, p]
```

```
{5/12, 1/4, 1/3}
```

"Summing over " is also called **marginalization** or "**integrating out**". Note that marginalization turns a probability function with higher degrees of freedom into one of lower degrees of freedom.

What are the prior probabilities? $p(H)$? To find out, we sum over the rows:

```
pH = Apply[Plus, Transpose[p]]
```

```
{1/3, 2/3}
```

■ Conditioning

Now that we have the marginals, we can get use the *product rule* to obtain the conditional probability through conditioning of the joint:

$$p(x|H) = \frac{p(H, x)}{p(H)}$$

In the Exercises, you can see how to use *Mathematica* to do the division for conditioning. The syntax is simple:

```
pXH = p / pH
```

```
{1/4, 1/4, 1/2}
{1/2, 1/4, 1/4}
```

Note that the probability of x conditional on H sums up to 1 over x , i.e. each row adds up to 1. But, the columns do not. $p(x|H)$ is a **probability** function of x , but a **likelihood** function of H . The posterior probability is obtained by conditioning on x :

$$p(H|x) = \frac{p(H, x)}{p(x)}$$

Syntax here is a bit more complicated, because the number of columns of p_X don't match the number of rows of p . We use `Transpose[]` to exchange the columns and rows of p before dividing, and then use `Transpose` again to get back the 2×3 form:


```
pHX = Transpose[Transpose[p] / pX]
```

```
{1/5, 1/3, 1/2}
{4/5, 2/3, 1/2}
```

Plotting the joint

The following `BarChart[]` graphics function requires in add-in package (`<< Graphics`Graphics``), which is specified at the top of the notebook. You could also use `ListDensityPlot[]`.

```
BarChart[p[[1]], p[[2]]];
```



Marginalization and conditioning: An example using *Mathematica* functions

■ A discrete joint probability

All of our knowledge regarding the signal discrimination problem can be described in terms of the joint probability of the hypotheses, H and the possible data measurements, x . The probability function assigns a number to all possible combinations:

$p(H, x)$

That is, we are assuming that both the hypotheses and the data are discrete random variables.

```
H = {S1, S2}
x ∈ {1, 2, ...}
```

Let's assume that x can only take on one of three values, 1, 2, or 3. And suppose the joint probability is:

```
p[H_, x_] := Which[H == 1 && x == 1, 1/12, H == 1 && x == 2, 1/12, H == 1 && x == 3,
1/6, H == 2 && x == 1, 1/3, H == 2 && x == 2, 1/6, H == 2 && x == 3, 1/6];
```

```
TableForm[Table[p[H, x], {H, 1, 2}, {x, 1, 3}],
TableHeadings -> {"H=s1", "H=s2"}, {"X=1", "X=2", "X=3"}]
```

	X=1	X=2	X=3
H=s1	$\frac{1}{12}$	$\frac{1}{12}$	$\frac{1}{6}$
H=s2	$\frac{1}{3}$	$\frac{1}{6}$	$\frac{1}{6}$

The total probability should sum up to one. Let's test to make sure. We first turn the list of lists into a single list of scalars using `Flatten[]`. And then we can sum either with `Apply[Plus,Flatten[p]]`.

```
Sum[p[H, x], {H, 1, 2}, {x, 1, 3}]
```

```
1
```

We can pull out the first row of `p` like this:

```
Table[p[1, x], {x, 1, 3}]
```

```
{ $\frac{1}{12}$ ,  $\frac{1}{12}$ ,  $\frac{1}{6}$ }
```

Is this the probability of `x`? No. For a start, the numbers don't sum to one. But we can get it through the two processes of marginalization and conditioning.

■ Marginalizing

What are the probabilities of the data, $p(x)$? To find out, we use the *sum rule* to sum over the columns:

```
px[x_] := Sum[p[H, x], {H, 1, 2}];
```

```
Table[px[x], {x, 1, 3}]
```

```
{ $\frac{5}{12}$ ,  $\frac{1}{4}$ ,  $\frac{1}{3}$ }
```

"Summing over" is also called **marginalization** or "**integrating out**". Note that marginalization turns a probability function with higher degrees of freedom into one of lower degrees of freedom.

What are the prior probabilities? $p(H)$? To find out, we sum over the rows:

```
pH[H_] := Sum[p[H, x], {x, 1, 3}];
```

```
Table[pH[H], {H, 1, 2}]
```

```
{ $\frac{1}{3}$ ,  $\frac{2}{3}$ }
```

■ Conditioning

Now that we have the marginals, we can get use the *product rule* to obtain the conditional probability through conditioning of the joint:

$$p\{x|H\} = \frac{p\{H, x\}}{p\{H\}}$$

We use function definition in *Mathematica* to do the division for conditioning. The syntax is simple:

```
pxH[H_, x_] := p[H, x] / pH[H];
```

```
Table[pxH[H, x], {H, 1, 2}, {x, 1, 3}]
```

```
{ $\frac{1}{4}$   $\frac{1}{4}$   $\frac{1}{2}$   
 $\frac{1}{2}$   $\frac{1}{4}$   $\frac{1}{4}$ }
```

Note that the probability of `x` conditional on `H` sums up to 1 over `x`, i.e. each row adds up to 1. But, the columns do not. $p\{x|H\}$ is a **probability** function of `x`, but a **likelihood** function of `H`. The posterior probability is obtained by conditioning on `x`:

$$p\{H|x\} = \frac{p\{H, x\}}{p\{x\}}$$

```
pHx[H_, x_] := p[H, x] / px[x];
```


```
Table[pHx[H, x], {H, 1, 2}, {x, 1, 3}]
```

```
{ $\frac{1}{5}$   $\frac{1}{3}$   $\frac{1}{2}$   
 $\frac{4}{5}$   $\frac{2}{3}$   $\frac{1}{2}$ }
```

Plotting the joint

We use `ListDensityPlot[]`.

```
ListDensityPlot[Table[p[H, x], {H, 1, 2}, {x, 1, 3}]];
```



Exercises

Exercise: Use density mapping theorem to make random number generator for density $p(y)$

References

- Applebaum, D. (1996). *Probability and Information*. Cambridge, UK: Cambridge University Press.
- Cover, T. M., & Joy, A. T. (1991). *Elements of Information Theory*. New York: John Wiley & Sons, Inc.
- Duda, R. O., & Hart, P. E. (1973). *Pattern classification and scene analysis*. New York.: John Wiley & Sons.
- Golden, R. (1988). A unified framework for connectionist systems. *Biological Cybernetics*, *59*, 109-120.
- Kersten, D. and P.W. Schrater (2000). *Pattern Inference Theory: A Probabilistic Approach to Vision*, in *Perception and the Physical World*, R. Mausfeld and D. Heyer, Editors. , John Wiley & Sons, Ltd.: Chichester. (pdf)
- Kersten, D., Mamassian P & Yuille A (in press) Object perception as Bayesian inference. *Annual Review of Psychology*. (pdf, <http://arjournals.annualreviews.org/doi/pdf/10.1146/annurev.psych.55.090902.142005>)
- Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology*, *13*(2) <http://gandalf.psych.umn.edu/~kersten/kersten-lab/papers/KerstenYuilleApr2003.pdf>
- Knill, D. C., & Richards, W. (1996). *Perception as Bayesian Inference*. Cambridge: Cambridge University Press.
- Ripley, B. D. (1996). *Pattern Recognition and Neural Networks*. Cambridge, UK: Cambridge University Press.
- Van Trees, H. L. (1968). *Detection, Estimation and Modulation Theory*. New York: John Wiley and Sons.
- Yuille, A., Coughlan J., Kersten D.(1998) (pdf)